


# Remote Imagery and Supercomputing:

## A Match Made “OnEarth”

Lucian Plesea


Jet Propulsion Laboratory




# Synopsis


---

- OnEarth is the name of the public JPL Web Map Service (WMS) server, hosting a half arc-second, global earth image mosaic.
- Has 3600 times more pixels than previous global earth images, at about 15m per pixel
- This mosaic contains about three trillion pixels, more than five Terabytes of data.
- Creating such a large mosaic was only possible by using large parallel computers.



# World Mosaic






# The Project

---

- NASA is leading an effort to increase the accessibility of remote imagery.
- The availability of a recent, high resolution, global coverage map of the earth was seen as an important component of this effort.
- The project started in earnest in Jan 2003, with the expected completion in Dec 2003.





# Level of Detail





# Custom Images






# Components

---


- Dataset: GeoCover, orthorectified Landsat 7
- Storage: RASCHAL, a 40TB storage cluster
- File Storage: Journaling Image File Format
- Application: Parallel image mosaic engine
- Data Access: Network Image Protocol
- Computation Resource: IPG
- Mosaic: Release 1
- Mosaic access: WMS



# Dataset

---

- NASA purchased global orthorectified Landsat 7 as a part of the Scientific Data Purchase Phase II.
- Expected 8000 scenes, with a final delivery date of July 2003. Currently 10600 scenes have been received.
- Each scene has 9 spectral bands, in UTM coordinates, and is about 500 Megabytes




# Storage

- For this project, and a few others, a very large storage system was necessary.
- Raid Again Storage using Commodity Hardware And Linux (Raschal), is a 40TB NAS storage cluster, built in-house.
- RASCHAL became operational April 2003, and has been in continuous use since then.
- OnEarth uses about 20TBytes

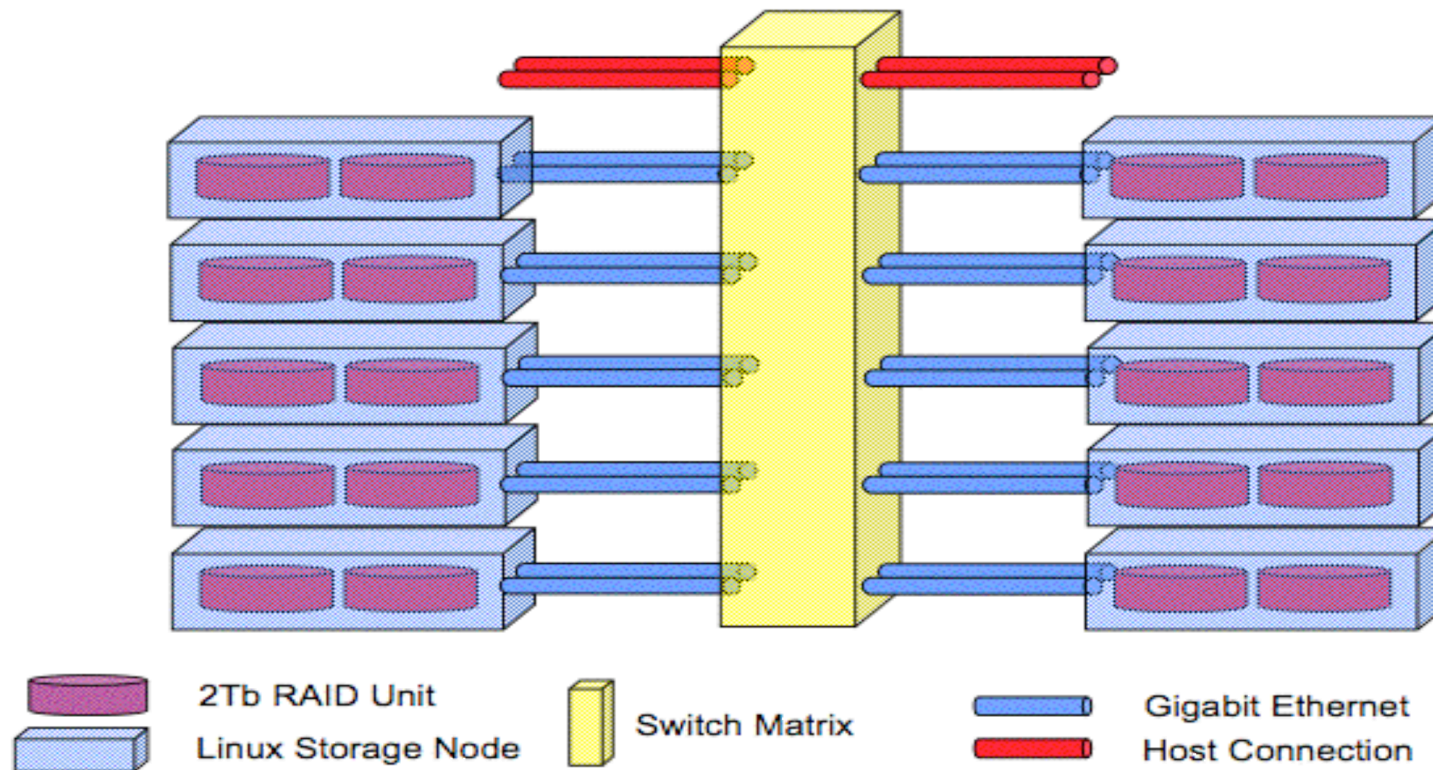







# Storage

RASCHAL System Diagram






# File Format

---

- A Journaling Image File Format is used extensively, for both the input scenes and the output mosaic.
- It is a tiled, multispectral and multi-resolution file format that supports lossless and lossy compression at the tile level.
- A level of indirection in data access, adding journaling features which ensures file consistency






# File Format

---


- Input images have been converted to this format, using bz2 (block arithmetic) lossless compression.
- Lower resolution input images are pre-built, for lower resolution test runs.
- The output mosaic itself is stored using libz compression, less efficient but faster.
- Files larger than 1Tbyte have been generated during this project



# Components

---


- Dataset: GeoCover, orthorectified Landsat 7
- Storage: RASCHAL, a 40TB storage cluster
- File Storage: Journaling Image File Format
- Application: Parallel image mosaic engine
- Data Access: Network Image Protocol
- Computation Resource: IPG
- Mosaic: Release 1
- Mosaic access: WMS



# Application

---


- An image mosaic builder for very large datasets, applying in a single pass the coordinate transformation, color correction and scene blending.
- Unique capabilities include UTM to geographical reprojection , blend mask driven data selection and feathering, and per band first order polynomial color correction



# Application


---


- In a first phase, the input scenes are analyzed, data location and the statistical distribution of each spectral band are determined. This data is saved for later use.
- The data collected in the first stage is then used to generate a global color-matching solution. Blend masks are generated, derived from the location information.




# Input Scene Example


---





# Blend Mask






# Application

---


- Many algorithms for the color-match have been tried. The current approach uses the MODIS derived, “Blue Marble” 30 arc-second image brightness to direct the Landsat 7 panchromatic band brightness. The other bands are using the coefficients from the panchromatic band, followed by a recursive filtering to reduce neighbor scene differences.






# Panchromatic Band

---






# Application

---


- The mosaic engine reads all the input bands required for a small area of the output.
- The optional color-matching correction is applied, then transforms the scene to the output projection space.
- The images are now in a common projection and can be blended together, a normalized weighted sum controlled by the blend masks.



# Application

---


- Since the mosaic engine only operates on a small area of output at a time, checkpoint and restart capabilities are built in.
- Use of the Journaling File Format allows for multiple processing cluster to generate output for the same file at the same time.
- The mosaic builder code is using a shared memory architecture, and uses SGI's Image Vision Library.



# Data Access

---


- Since the computation resources are located remote from the storage resources, direct access to the data is not possible.
- An image specific data access subsystem allows small regions of the input and output images to be transferred independently. The Network Image Protocol is used, separating the location and specific file format from the application



# Data Access

---


- Tiles can be transferred in either raw or compressed format, under user control.
- A single plain text setup file provides access to all input files, in this case more than 70,000
- Since the Network Image Protocol is available as an Image Library module, no changes were required to the mosaic engine



# Computation Resource

---

- Since the mosaic engine is a shared memory application, it benefits from faster CPUs and a large memory footprint.
- A 1024 CPU Origin 3000 machine, with 600MHz CPUs and 256MB per CPU was used.
- This machine is located at AMES, and connected to JPL via a 15MB/sec link, the NREN network component of the IPG




# Computation Resource

---

- Four groups of 32 CPUs were running simultaneously to produce the mosaic, using about 50000 CPU-hours.
- This architecture provided a balance of data access and computational loads, achieving a CPU load between 30 and 60 percent.
- Peak data transfer rates of 10MB/sec were seen, with an average of 2.5MB/sec data read and 1.25MB/sec data write






# Components

---


- Dataset: GeoCover, orthorectified Landsat 7
- Storage: RASCHAL, a 40TB storage cluster
- File Storage: Journaling Image File Format
- Application: Parallel image mosaic engine
- Data Access: Network Image Protocol
- Computation Resource: IPG
- Mosaic: Release 1
- Mosaic access: WMS



# WMS Global Mosaic

---


- There are nine separate images, at three different resolutions. Coverage area is 180W to 180E and S85 to N85.
- The panchromatic band, at 0.5 arc-second or 15m, a 2,592,000x1,224,000 image.
- The thermal band, with high and low gains, at 2 arc-second per pixel (60m).
- The three visual and three near IR bands, stored at 1 arc-second (30m).



# WMS Global Mosaic

---


- Each resolution is stored in a separate file.
- The base resolution, in the native Journaling File Format, takes 1.3TeraBytes of storage.
- The pan band is about 500GB, the thermal bands 35GB and the normal bands 800GB
- Adding multiple resolutions (13 levels), will bring the dataset to about 1.8TeraBytes




# Access: WMS

---

- Access to the mosaic is best done via the WMS server. A prototype site is available on <http://OnEarth.jpl.nasa.gov>
- The server is implemented as a CGI application, and uses the same technologies.
- Provides, color selection, pan-sharpening, multiple projections and image control using Styled Layer Descriptor.





# To Be Continued


---

- Pending:
  - WMS access to mosaic
  - Web Site Refresh
  - Packaging for delivery to interested parties
- Possible Further Development
  - Improve and maintain the mosaic
  - SRTM Elevation
  - Web Terrain Service










# Credits

---

- Dave Curkendall, shares blame for the whole thing
- Kacie Shelton, the tape and data shepherd
- Jimi Patel, built most of RASCHAL
- Richard Schreyer, SURF student, wrote the current WMS engine.
- George Percivall, manages the WMS Global Mosaic project for NASA



# Sponsors

---

- Sponsors:
  - WMS Global Mosaic
    - Geospatial Interoperability Office, YO
  - Image Access Technologies:
    - ESTO-CT
  - Computing Resources:
    - AMES CNIS Program